

CLAIMS

WE CLAIM:

1. A method of coordinating at least two processor units, each having a processor and cache memory, and communicating cache coherence messages with each other and a shared memory over a network, the method comprising the steps of:

(a) providing a mechanism for communications of cache coherence messages

5 directly from a given processor unit to another processor unit;

(b) providing a mechanism for communication of cache coherence messages directly from a given processor unit to a directory and then to at least one other processor unit when indicated by the directory;

(c) evaluating the available bandwidth on the network used to communicate
10 the cache coherence messages; and

(d) for a given cache coherence message, selecting one the mechanism of step (a) or the mechanism of step (b) based on the evaluation of step (c).

2. The method recited in claim 1 wherein the mechanism of step (a) broadcasts the cache coherence message to all other processor units.

3. The method recited in claim 1 wherein the given cache coherence message identifies a block of the shared memory and wherein the directory provides an index linking blocks of memory to a set of processor units less than all the processor units and wherein the mechanism of step (b) sends the cache coherence
5 message to the given set of processor units linked to the block of shared memory identified by the given cache coherence message.

4. The method recited in claim 3 wherein the directory sends the cache coherence message directly over the network to the given set of processor units.

5. The method recited in claim 3 including the step of:

(e) detecting insufficiency in the set of processor units to which the coherence message is transmitted;

5 (f) retrying the transmission of the cache coherence message to the given set of processor units; and

(g) upon repeated insufficiency in the set of processor units to which the transmission of the coherence message is retried in step (f), broadcasting the given cache coherence message to all processor units.

6. The method of claim 5 wherein retries of the transmission of the cache coherence message append a retry number to the cache coherence message and responses to the cache coherence message.

7. The method recited in claim 5 wherein the number of retries is limited to predetermined number less than ten.

8. The method recited in claim 1 wherein the evaluation of available bandwidth compares the available bandwidth against a predetermined threshold and selects the mechanism of step (a) when the available bandwidth is greater than the threshold and selects the mechanism of step (b) when the available bandwidth is less
5 than the threshold.

9. The method recited in claim 8 wherein the threshold is less than all of the bandwidth of the network.

10. The method recited in claim 9 wherein the threshold is substantially 75% of the capacity of the network.

11. The method recited in claim 1 wherein step (d) provides for successive given cache coherence messages being transmitted using a mix of selections of the mechanism of step (a) and the mechanism of step (b) the mix being a function of the evaluation of step (c) to provide a semicontinuous variation in the mix.

12. The method recited in claim 11 wherein the selection of the mechanism of step (a) or the mechanism of step (b) for a given cache coherence message is done pseudorandomly according to a probability function based on the evaluation of step (c).

13. The method recited in claim 1 wherein the mechanism of step (a) multicasts the cache coherence message to a selected set of processor units based on a prediction as to which processor units have cache memories loaded with relevant data.

14. The method recited in claim 13 wherein a directory monitors the multicasting to detect insufficiency in the targets of the multicasting resulting from erroneous prediction and to initiates a retransmission of the cache coherence message.

15. The method recited in claim 1 wherein step (c) of evaluating the available bandwidth monitors the communications on the network at the given processor unit transmitting the given cache coherence message.

16. The method recited in claim 1 wherein the mechanism of step (b) communicates cache coherence messages directly from a given processor unit to a directory and to the given processor unit.

17. A method of coordinating at least two processor units, each having a processor and cache memory, and communicating cache coherence messages with each other and a directory over a network, the method comprising the steps of:

(a) multicasting from a given processor unit, a cache coherence message to a selected set of other processor units, based on a prediction as to which other processor units have cache memories loaded with relevant data;

(b) using the directory to detect insufficiency in the selected set of other processor units to which transmission of the cache coherence message is made; and

(c) upon a detected insufficiency, causing the directory to retry the multicast transmission of the cache coherence message.

18. The method recited in claim 17 including the step of
(d) upon repeated insufficiency in step (c), broadcasting the given cache coherence message to all processor units.

19. The method recited in claim 17 wherein the repeated insufficiency is a predetermined number less than ten.

20. The method recited in claim 17 wherein the directory sends the retry multicast transmissions to processor units likely to have the relevant data based on a monitoring of cache coherence messages from processor units.

21. The method of claim 17 wherein the directory appends a retry number to retries of the cache coherence message.

22. The method of claim 21 wherein the processor units responding to the retries appends the retry number to the responses to the retried cache coherence message.

23. The method recited in claim 17 wherein at step (c) the multicast transmission of the cache coherence message is also sent to the given processor unit originating the cache coherence message.

24. Cache-coherence circuitry for a computer architecture having: (a) a shared memory, (b) at least two processor units, each having a processor and cache memory, and (c) a network for communicating cache coherence messages among the processor units and the shared memory, the cache-coherence circuitry

5 comprising:

(a) snooping means for communications of cache coherence messages directly from a given processor unit to another processor unit;

(b) directory means for communication of cache coherence messages directly from a given processor unit to a directory and then to at least one other processor
10 unit when indicated by the directory;

(c) evaluation means for evaluating the available bandwidth on the network used to communicate the cache coherence messages; and

(d) selection means for choosing one the snooping means and directory means for the communication of a given cache coherence message based on the evaluation of available bandwidth determined by the evaluation means.

25. The cache coherence circuitry recited in claim 24 wherein the snooping means broadcasts the cache coherence message to all other processor units.

26. The cache coherence circuitry recited in claim 24 wherein the given cache coherence message identifies a block of the shared memory and wherein the directory provides an index linking blocks of memory to a set of processor units less than all the processor units and wherein the directory means sends the cache coherence message to the given set of processor units linked to the block of shared memory identified by the given cache coherence message.

27. The cache coherence circuitry recited in claim 26 wherein the directory means sends the cache coherence message directly over the network to the given set of processor units.

28. The cache coherence circuitry recited in claim 26 further including monitoring means for:

(i) detecting insufficiency in the set of processor units to which the coherence message is transmitted;

5 (ii) retrying the transmission of the cache coherence message to the given set of processor units; and

(iii) upon repeated insufficiency in the set of processor units to which the transmission of the coherence message is retried, broadcasting the given cache coherence message to all processor units.

29. The cache coherence circuitry recited in claim 28 wherein the monitoring means appends a retry number to the cache coherence message and responses to the cache coherence message.

30. The cache coherence circuitry recited in claim 28 wherein the number of retries is limited to predetermined number less than ten.

31. The cache coherence circuitry recited in claim 24 wherein the evaluation means compares the available bandwidth against a predetermined threshold and selects the snooping means when the available bandwidth is greater than the threshold and selects the directory means when the available bandwidth is less than
5 the threshold.

32. The cache coherence circuitry recited in claim 31 wherein the threshold is less than all of the bandwidth of the network.

33. The cache coherence circuitry recited in claim 32 wherein the threshold is substantially 75% of the capacity of the network.

34. The cache coherence circuitry recited in claim 24 wherein the selection means provides for successive given cache coherence messages being transmitted using a mix of the snooping means and the directory means the mix being a function

of the evaluation of available bandwidth by the evaluation means to provide a
5 semicontinuous variation in the mix.

35. The cache coherence circuitry recited in claim 34 wherein the selection of the snooping means or the directory means for a given cache coherence message is done pseudorandomly according to a probability function based on the evaluation of available bandwidth by the evaluation means.

36. The cache coherence circuitry recited in claim 24 wherein the snooping means multicasts the cache coherence message to a selected set of processor units based on a prediction as to which processor units have cache memories loaded with relevant data.

37. The cache coherence circuitry recited in claim 36 including a monitoring means monitoring the multicasting to detect insufficiency in the targets of the multicasting resulting from erroneous prediction and to initiates a retransmission of the cache coherence message.

38. The cache coherence circuitry recited in claim 24 wherein the evaluation of available bandwidth by the evaluation means monitors the communications on the network at the given processor unit transmitting the given cache coherence message.

39. The cache coherence circuitry recited in claim 24 wherein the directory means communicates cache coherence messages directly from a given processor unit to a directory and to the given processor unit.

40. Cache-coherence circuitry for a computer architecture having: (a) a shared memory, (b) at least two processor units, each having a processor and cache memory, and (c) a network for communicating cache coherence messages among the processor units and the shared memory, the cache-coherence circuitry
5 comprising:

(a) predictive multicasting circuitry, multicasting from a given processor unit, a cache coherence message to a selected set of other processor units, based on a

prediction as to which other processor units have cache memories loaded with relevant data; and

- 10 (b) a directory detecting insufficiency in the selected set of other processor units to which transmission of the cache coherence message is made, the directory operating upon a detected insufficiency, to retry the multicast transmission of the cache coherence message.

41. The cache coherence circuitry recited in claim 40 wherein the directory further, upon repeated insufficiency in the selected set of other processor units, broadcasts the given cache coherence message to all processor units.

42. The cache coherence circuitry recited in claim 40 wherein the repeated insufficiency is a predetermined number less than ten.

43. The cache coherence circuitry recited in claim 40 wherein the directory sends the retry multicast transmissions to processor units likely to have the relevant data based on a monitoring of cache coherence messages from processor units.

44. The method of claim 40 wherein the directory appends a retry number to retries of the cache coherence message.

45. The method of claim 44 including circuitry within the processor units responding to the retries appends the retry number to the responses to the retried cache coherence message.

46. The cache coherence circuitry recited in claim 40 wherein the predictive multicasting circuitry sends the cache coherence message also to the given processor unit originating the cache coherence message.